II. REMARKS

Independent claims 1 and 9 have been amended to additionally recite that "said processor of each memory module manages said portion of series of data such that said series of data is divided among the plurality of memory modules" as supported on page 4, lines 8-25, and on page 10, lines 8-11 and lines 22-29, of the specification as originally filed.

No new matter has been added by the above amendments.

A. The Invention

The presently claimed invention pertains broadly to apparatuses having parallel computing architecture, such as is used to implement a Single Instruction Stream, Multiple Data Stream ("SIMD") architecture for performing general purpose parallel processing using appropriate and high-speed memory control. More specifically, the present invention pertains to a computer system having architecture of a parallel computer, wherein the computer system has the features recited in claim 1 of the present application. In addition, the present invention also pertains to an information processing unit having the features recited in claim 9 of the present application.

Various other embodiments, in accordance with the present invention, are recited in the dependent claims. All of the various embodiments, in accordance with the present invention, advantageously utilize an architecture of a parallel computer so that parallel processing by means of appropriate and high-speed memory control can be achieved.

B. The Rejections

Claims 1-13 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Dowling (U.S. Patent 6,226,738 B1, hereafter the "Dowling Patent") in view of Miller et al. (U.S. Patent 5,490,260, hereafter the "Miller Patent") and Chang et al. (U.S. Reissue Patent RE 37,305 E, hereafter, the "Chang Patent").

Applicant respectfully traverses the rejection and requests reconsideration of the instant claims for the following reasons.

C. Applicant's Arguments

As best as Applicant can understand the outstanding March 9th Office Action, the Examiner contends the combination of the Dowling Patent, the Miller Patent and the Chang Patent would render obvious the "processing of the accessed data" as recited by independent claims 1 and 9 of the above-captioned application because "[o]ne of ordinary skill would have been motivated [to] update the address table when data was stored at least to provide reliable access to data" because "without updated address table the system would be vulnerable to inadvertently overwriting data already stored in the memory modules" (Office Action, dated March 9, 2006, at 5, lines 11-16).

In the present case, the Examiner's Section 103 rejection, which is based on the combined teachings of the Dowling Patent, the Miller Patent and the Chang Patent, is facially flawed because the Examiner attempts to employ the doctrine of inherency to make up deficiencies in the disclosure of the combined patents. However, the doctrine of inherency does not apply to Section 103 rejections involving combinations of multiple references. See Continental Can Co. v. Monsanto Co., 20 U.S.P.Q.2d 1746, 1749 (Fed. Cir. 1991).The doctrine of inherency cannot be used to infer "inherent"

features allegedly resulting when teachings are combined. To the extent that the Examiner improperly applies the doctrine of inherency in making the present Section 103 rejection, Applicant traverses the rejection and asserts that it is, therefore, facially untenable and must be withdrawn. To the extent that the Examiner is taking notice of the allegedly inherent features missing from the references, this notice is traversed.

D. The Section 103 Rejection

A prima facie case of obviousness requires a showing that the scope and content of the prior art teaches each and every element of the claimed invention, and that the prior art provides some teaching, suggestion or motivation to combine the references to produce the claimed invention. In re Oetiker, 24 U.S.P.Q.2d 1443 (Fed. Cir. 1992); In re Vaeck, 20 U.S.P.Q.2d 1438 (Fed. Cir. 1991).

In the present case, the Examiner's rejection of the instant claims 1-13 under 35 U.S.C. § 103 is clearly untenable because the prior art fails to teach all of the claimed subject matter.

i. The Dowling Patent

The Dowling Patent teaches a "split embedded dram processor," such as shown in Figures 1-3, wherein a CPU core (100), (200), or (300), is coupled via a standard memory bus connection (120) to a plurality of memory modules (140) contained in an embedded DRAM (110), (210), or (310). The Dowling Patent teaches that the embedded DRAM (110) includes a DRAM memory array (140) coupled to an embedded logic CPU extension (150) via internal bussing structure (160), (col. 13, lines 35-64). As shown in Figure 4, the Dowling Patent further teaches that the embedded DRAM has a program space (410) that includes type II instructions T2'(j)

and type III instructions T3(k), (col. 16, lines 15-17). However, as conceded by the Examiner, the Dowling Patent does not teach, or suggest, any detail regarding table management (Office Action, dated March 9, 2006, at 3, line 1). Thus, the Dowling Patent does not teach, or suggest, that

> "architecture of a parallel computer manages at least one series of data having a stipulated relationship, each series of data being given a space ID, and the processor of each memory module manages a table that contains one or more sets of said space ID, the logical address of a portion of the series of data managed, the size of said portion and the size of the series of data, in which said processor of each memory module manages said portion of series of data such that said series of data is divided among the plurality of memory modules and, in response to an instruction including the space ID and the logical address from the CPU module, the processor of each memory module determines if the portion of the series of data managed is involved in the received instruction by reviewing the space ID and the logical address, reads data stored in the RAM core and sends data out on a bus, writes data given via the bus to the RAM core, performs (c) the necessary processing on the data or (d) updates said table, or performs (c) and (d)"

as recited in independent claims 1 and 9 of the present application.

The Dowling Patent simply does not teach, or suggest, that the processor (i.e., CPU extension logic) of each memory module (embedded DRAM) manages a portion of the series of data such that the series of data, having the same space ID, is divided among the plurality of memory modules (embedded DRAMs). Instead, the Dowling Patent merely teaches that the memory array in the embedded DRAM has program space and stores the type II and type III instructions.

According to the presently claimed invention, the "series of data is divided among the plurality of memory modules and, in response to an instruction...the processor of each memory module determines if the portion of the series of data managed is involved in the received instruction...." Thus, in accordance with the presently claimed invention, while the "series of data is divided among the plurality of

memory modules," it is not necessary that the series of data be divided equally into the number of memory modules of the plurality and it is not necessary that every memory module of the plurality of memory modules be involved in storing and managing the divided series of data. In other words, a person of ordinary skill in the art would recognize that it falls within the scope of the claimed invention for the series of data to be divided into some number portion of the plurality of memory modules, and each of the involved memory modules stores and manages a portion of the divided series of data. According to claim 1, as amended, the processor of each memory module, of those memory modules involved in storing and managing the divided series of data, manages a portion of the divided series of data such that the series of data is divided among the plurality of memory modules.

ii. The Miller Patent

The Miller Patent teaches a "solid-state RAM data storage for virtual memory computer using fixed-sized swap pages with selective compressed/uncompressed data store according to each data size," which uses virtual memory management that employs a random-access memory type storage device (i.e., semiconductor memory) for page swapping (See Abstract). The semiconductor memory is formatted to provide multiple partitions of varying block size for compressed pages and uncompressed pages; however, the data is stored in pages of fixed size and these pages are compressed for storage if the compressed size fits in the block size of one of the small-block partitions in the memory (See Abstract). If a data page is not compressible to one of the small block sizes, it is stored uncompressed in the other full-sized partition (See Abstract). As conceded by the Examiner, the Miller Patent does not teach, or suggest, a

"space ID" as recited in claims 1 and 9 (Office Action, dated March 9, 2006, at 4, line 6).

However, the Miller Patent also does not teach, or even suggest,

> "each series of data being given a space ID, and the processor of each memory module manages a table that contains one or more sets of said space ID, the logical address of a portion of the series of data managed, the size of said portion and the size of the series of data, in which said processor of each memory module manages said portion of series of data such that said series of data is divided among the plurality of memory modules"

as recited in claims 1 and 9. In other words, the Miller Patent does not teach, or suggest, that each memory module manages a portion of the divided series of data such that the series of data having the same space ID is divided up among the plurality of memory modules as discussed above..

iii. The Chang Patent

The Chang Patent teaches a "virtual memory address translation mechanism with controlled data persistence" that employs a segment table format, as shown in Figure 2, wherein a 12-bit segment identifier is concatenated with bits 4 through 31 of an incoming effective address to form a 40-bit virtual address (col. 11, lines 4-6). The lower order 11 bits for 2K pages, or 12 bits for 4K pages, of the effective address are used as the byte address for the selected real page, and these bits are not altered by the translation process (col. 11, lines 6-9). The remaining 29(28) bits of the virtual address are then presented to the translation hardware and the generation of the virtual address uses the segment identifier and the storage effective address (col. 11, ones 9-13).

However, the Chang Patent does not teach, or even suggest,

> "each series of data being given a space ID, and the processor of each memory module manages a table that contains one or more sets of said space ID, the logical address of a portion of the series of data managed, the size of said portion and the size of the series of data, in which said processor of each

> memory module manages said portion of series of data such that said series of data is divided among the plurality of memory modules"

as recited in claims 1 and 9. In other words, the Chang Patent does not teach, or suggest, that each memory module manages a portion of the divided series of data such that the series of data having the same space ID is divided up among the plurality of memory modules as discussed above.

iv. Summary of the Dowling, Miller and Chang Patents

Neither the Dowling Patent, the Miller Patent, nor the Chang Patent teaches, or suggests,

> "each series of data being given a space ID, and the processor of each memory module manages a table that contains one or more sets of said space ID, the logical address of a portion of the series of data managed, the size of said portion and the size of the series of data, in which said processor of each memory module manages said portion of series of data such that said series of data is divided among the plurality of memory modules"

as recited in independent claims 1 and 9; therefore, no combination of these patents can teach this feature of the claimed invention. For this reason, the Examiner has failed to establish a prima facie case of obviousness against claims 1 and 9 of the present application.

v. Claim 7

The Examiner concedes that neither the Dowling Patent, the Miller Patent, nor the Chang Patent teach, or suggest

> "wherein, when said processor receives an instruction to delete a specific element within a series of data, insert a specific element into said series of data, or add a specific element to the end of a series of data, said processor performs a table lookup, compares the region of data managed against the position of said element subject to deletion, insertion or addition, and based on the results of said comparison, updates the content of said table"

as recited by claim 7 (Office Action, dated March 9, 2006, at 6, lines 9-17). Instead, the Examiner contends that such operations as deletion of specific elements, insertion of specific elements, or adding a specific element to the end of a series of data, are well known in the art and would be "anticipated" by the combined teachings of the Dowling, Miller and Chang Patents (Office Action, dated March 9, 2006, at 6, lines 9-17).

Applicant respectfully objects to the Examiner's contention regarding what the Examiner believes is well known in the art and requests that the Examiner substantiate his position with a prior art reference. It is a well-settled proposition that an Examiner's rejection must be based on evidence, and not on subjective belief and unknown authority. In re Lee, 61 U.S.P.Q.2d 1430, 1434 (Fed. Cir. 2002).

The Examiner must now produce evidence, such as a prior art reference, to support the Examiner's contention regarding what the Examiner alleges is the well-known subject matter or withdraw the rejection against claim 7.

III. CONCLUSION

The rejection under 35 U.S.C. § 103(a) is untenable and should be withdrawn because neither the Dowling Patent, the Miller Patent nor the Chang Patent reasonably teaches, or even suggests, "each series of data being given a space ID, and the processor of each memory module manages a table that contains one or more sets of said space ID, the logical address of a portion of the series of data managed, the size of said portion and the size of the series of data, in which said processor of each memory module manages said portion of series of data such that said series of data is divided among the plurality of memory modules" as recited in independent claims 1 and 9.

For all of the above reasons, claims 1-13 are in condition for allowance and a prompt Notice of Allowance is earnestly solicited.

Questions are welcomed by the below-signed attorney for applicant.

Respectfully submitted,

GRIFFIN & SZIPL, PC

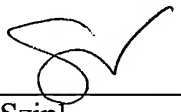

Joerg-Uwe Szipl
Registration No. 31,799

GRIFFIN & SZIPL, PC
Suite PH-1
2300 Ninth Street, South
Arlington, VA 22204
Telephone: (703) 979-5700
E-mail: gands@szipl.com
Facsimile: (703) 979-7429
Customer No.: 24203